

Oracle lower bounds for stochastic gradient sampling algorithms

NILADRI S. CHATTERJI^{1,*} PETER L. BARTLETT^{2,**} and PHILIP M. LONG^{3,†}

¹*Department of Physics, University of California Berkeley. E-mail: *chatterji@berkeley.edu*

²*Department of Statistics and EECS, University of California Berkeley. E-mail: **peter@berkeley.edu*

³*Google. E-mail: †plong@google.com*

We consider the problem of sampling from a strongly log-concave density in \mathbb{R}^d , and prove an information theoretic *lower bound* on the number of stochastic gradient queries of the log density needed. Several popular sampling algorithms (including many Markov chain Monte Carlo methods) operate by using stochastic gradients of the log density to generate a sample; our results establish an information theoretic limit for all these algorithms.

We show that for every algorithm, there exists a well-conditioned strongly log-concave target density for which the distribution of points generated by the algorithm would be at least ε away from the target in total variation distance if the number of gradient queries is less than $\Omega(\sigma^2 d / \varepsilon^2)$, where $\sigma^2 d$ is the variance of the stochastic gradient. Our lower bound follows by combining the ideas of Le Cam deficiency routinely used in the comparison of statistical experiments along with standard information theoretic tools used in lower bounding Bayes risk functions. To the best of our knowledge our results provide the first nontrivial dimension-dependent lower bound for this problem.

Keywords: Information theoretic lower bounds, Sampling lower bounds, Markov chain Monte Carlo, Stochastic gradient Monte Carlo.

1. Introduction

Sampling from a distribution is a crucial computational step in several domains such as Bayesian inference, prediction in adversarial environments, Monte Carlo approximation and reinforcement learning. In the high dimensional setting, a common approach is to use Markov chain Monte Carlo (MCMC) methods to generate a sample. Among these methods, particularly useful and widely applicable are MCMC algorithms that generate a sample from a probability distribution with a density on a continuous state space [41]. A broad class of such algorithms are first-order (gradient) methods. These methods operate by using the gradient of the log density at different points in the space.

Recently, the focus has been on methods that work with *stochastic* estimates of the gradient given the emergence of large-scale datasets [49]. On these datasets it is computationally challenging to calculate the exact gradient over the entire data. An example of

this, of course, is the case of sampling from a posterior distribution when the log density is sum decomposable over a large dataset.

While there is a large body of work studying MCMC methods in this setting (refer to [6] for a survey of the results in this area), classically, little was known about the explicit, *non-asymptotic* dependence of the iteration complexity—the number of steps of the method needed to get a good sample—of these sampling algorithms in terms of the dimension and the target accuracy. The seminal work of [11] established a non-asymptotic bound on the iteration complexity for the Unadjusted Langevin algorithm (ULA) [18, 36, 20, 43] in terms of the dimension and the target accuracy when the log density is sufficiently regular. This has kicked off a slew of research in this field and we now have many upper bounds on the non-asymptotic iteration complexity of several popular first-order sampling methods [14, 12, 9, 10, 15, 31, 32] in terms of relevant problem parameters like dimension, condition number and target accuracy.

In spite of this large and growing body of work, what has remained crucially missing is a characterization of the intrinsic hardness of sampling for a given family of distributions, that is, an information theoretic lower bound on the iteration complexity in terms of problem parameters for first-order sampling methods. This is the question that we address in our work.

We establish a lower bound on the iteration complexity for sampling from continuous distributions over \mathbb{R}^d that have smooth and strongly concave (see Section 2 for definitions) log densities of the form:

$$p^*(x) = \frac{\exp(-f(x))}{\int_{y \in \mathbb{R}^d} \exp(-f(y)) dy}, \quad \text{for all } x \in \mathbb{R}^d.$$

Prototypical examples of such distributions include the Gaussian distribution and the posterior distribution that arises in Bayesian logistic regression with a Gaussian prior.

To establish our lower bounds we borrow the *noisy oracle model* [35] used in the allied field of optimization. There, this model was used to establish lower bounds on the iteration complexity of optimizing a function using gradients/function-values [see also 39, 1, 40]. In this oracle model an algorithm (in our case a sampling method) is allowed to query the oracle at a point and the oracle returns the gradient/function value of the log density at that point corrupted by some independent and unbiased noise with bounded variance. The iteration complexity of the method is therefore equal to the number of queries that need to be made to this noisy oracle. The method can be adaptive, in the sense that it is allowed to choose its next query based on its entire history—past query points and the values returned by the oracle. Many popular first-order and zeroth-order (methods that work with function-values) sampling methods, like ULA, Hamiltonian Monte Carlo [34], Metropolis adjusted Langevin algorithm [43, 42, 5], Ball walk [27, 16, 28, 30], Metropolized random walk [33, 44], Hit-and-run [45, 2, 23, 26, 29, 30], underdamped Langevin MCMC [17, 10], Riemannian MALA [50], Proximal-MALA [37] and Projected ULA [7], can be described by this oracle model.

Our contributions. The main result of this paper is to establish a lower bound on the iteration complexity of any algorithm with access to a stochastic gradient oracle. We

show that the iteration complexity must grow at least linearly with the noise variance in the gradients.

More concretely, a sampling algorithm Alg operates in a sequential fashion. It begins by querying the stochastic gradient oracle at an initial point y_0 (which could be chosen in a randomized fashion). The oracle then returns the gradient of the negative log density at that point with some noise added to it: $z_0 = \nabla f(y_0) + \xi_0$, where ξ_0 is independent of the query point, has zero mean and has bounded variance $\mathbb{E}[\|\xi_0\|_2^2] \leq \sigma^2 d$. Such an assumption on the gradient oracle is quite common in the study of upper bounds of sampling algorithms (see, e.g [12, 10]). Denote the distribution of the noisy gradients as Q_x —the conditional distribution of the noisy gradient at a point x . **On observing this noisy gradient z_0 , the algorithm Alg can now choose to query at a new point y_1 , where this point y_1 is dependent on the initial query point y_0 and the noisy gradient z_0 .** This protocol continues for n rounds, after which the algorithm is asked to output a point from a distribution close to the target p^* . Let us denote the distribution of the point output by such an algorithm after n such queries by $\text{Alg}[n; Q]$.

We show that for every such algorithm, there exists a log smooth and strongly log-concave target density for which the distribution of the sample generated by the algorithm will be at least ε away from the target in total variation distance if the number of queries is less than $\Omega(\sigma^2 d / \varepsilon^2)$. A precise statement of this result **with explicit constants** is presented as Theorem 4.1.

This lower bound requires the choice of a different family of distributions for every target accuracy ε . A result that is easier to interpret is our lower bound for the case when the target accuracy is a constant. In this regime our results are as follows:

Corollary 1. *There exist constants c, c', C and C' such that for all d and σ^2 , $n < c\sigma^2 d$ implies*

$$\inf_{\text{Alg}} \sup_Q \sup_{p^*} \text{TV}(\text{Alg}[n; Q], p^*) > c',$$

where the infimum is over all gradient-based sampling algorithms, the supremum over Q is over unbiased gradient oracles with noise-variance bounded by $\sigma^2 d$, and the supremum over p^* is over C -log smooth, $C/2$ -strongly log-concave distributions over \mathbb{R}^d with the mean $\|\mathbb{E}_{\mathbf{x} \sim p^*} [\mathbf{x}]\|_2 \leq C'$.

This corollary asserts that to get to a distribution that is a constant away from a well-conditioned target distribution with constant log-smoothness it must make order $\sigma^2 d$ queries to the stochastic gradient oracle in the worst case.

We consider target distributions that have their mean in a ball of constant radius around the origin. This ensures that most of the mass of these distributions is concentrated in a ball of radius $\mathcal{O}(\sqrt{d})$ around the mean (as p^* is $C/2$ -strongly log-concave and hence $\sqrt{2/C}$ -sub-Gaussian).

At a high level the approach here is inspired by the work that proved lower bounds in stochastic optimization. There, the problem of lower bounding the error of any first-order optimization algorithm is reduced to that of lower bounding the Bayes risk of an

appropriately constructed statistical problem. We find that this reduction is not directly possible in our problem. Instead, we relate the problem of lower bounding the iteration complexity of sampling to that of lower bounding a *Le Cam deficiency* between two statistical experiments. Le Cam deficiency is a measure that is widely studied in the literature on the comparison of statistical experiments [4, 3, 24, 47] and it turns out to be a useful idea in our problem.

The rest of the paper is organized as follows. In Section 2, we introduce and define useful quantities used throughout the paper in the statement of our results and their proofs. In Section 3, we reduce the problem of lower bounding the iteration complexity of sampling algorithms to that of lower bounding a difference of Bayes risk functions. In Section 4, we state and prove our main results concerning gradient oracles. We conclude with a discussion in Section 5. The more technical details of the proofs are deferred to the appendix.

Notation. We use bold font to denote a random variable, and unbolded font to denote the particular sample value of that random variable. For example, \mathbf{x} denotes a random variable and $\mathbf{x} = x$ denotes that the realization of \mathbf{x} is x . We use uppercase letters to denote matrices and lower case letters to denote vectors. Given any vector v , let $\|v\|_2$ denote its Euclidean norm. Let $I_{d \times d}$ denote the identity matrix in d dimensions, $0_{d \times d}$ denote the matrix of all zeros and 0_d denote a row vector of all zeros. Given any subset $S \subseteq \mathbb{R}^d$, let $\mathcal{B}(S)$ denote the Borel σ -field associated with this set. For any finite set S , let $|S|$ denote its cardinality. $\mathbb{I}(\cdot)$ denotes the indicator of an event. We use $c, c', C, c_1, c_2, \dots$ to denote positive universal constants that are independent of any problem parameters. We adopt conventional notations $\mathcal{O}(\cdot)$ and $\Omega(\cdot)$ as suppressing constants independent of dimension and problem parameters, and $\tilde{\mathcal{O}}(\cdot)$ and $\tilde{\Omega}(\cdot)$ as suppressing poly-logarithmic factors.

2. Definitions and problem set-up

In this section, we define several quantities that are useful in stating and proving our results.

In our paper, we refer to f , the negative log density of the distribution p^* , as the potential function. As stated in the introduction, we consider a family of distributions with smooth and strongly convex potentials. We define smoothness and strong convexity below.

Definition 1. A function f from $\mathbb{R}^d \mapsto \mathbb{R}$ is α -smooth and β -strongly convex if there exists constants $\alpha \geq \beta > 0$ such that for all $x, y \in \mathbb{R}^d$,

$$\frac{\beta}{2} \|x - y\|_2^2 \leq f(x) - f(y) - \langle \nabla f(y), x - y \rangle \leq \frac{\alpha}{2} \|x - y\|_2^2.$$

If a function f is smooth and strongly convex then $-f$ is smooth and strongly concave. Next let us define a probability kernel.

Definition 2. For measurable spaces $(\mathcal{X}, \mathcal{S})$ and $(\mathcal{Y}, \mathcal{T})$, K is a probability kernel (conditional distribution) from \mathcal{X} to probability measures over \mathcal{Y} if

1. There is probability measure on $(\mathcal{Y}, \mathcal{T})$ associated with each $x \in \mathcal{X}$, which we denote by $K(\cdot|x)$.
2. The map $x \mapsto K(A|x)$ is \mathcal{S} -measurable for each $A \in \mathcal{T}$.

We will use $\mathbf{y} \sim K_x$ to represent the random variable \mathbf{y} drawn from the distribution $K(\cdot|x)$.

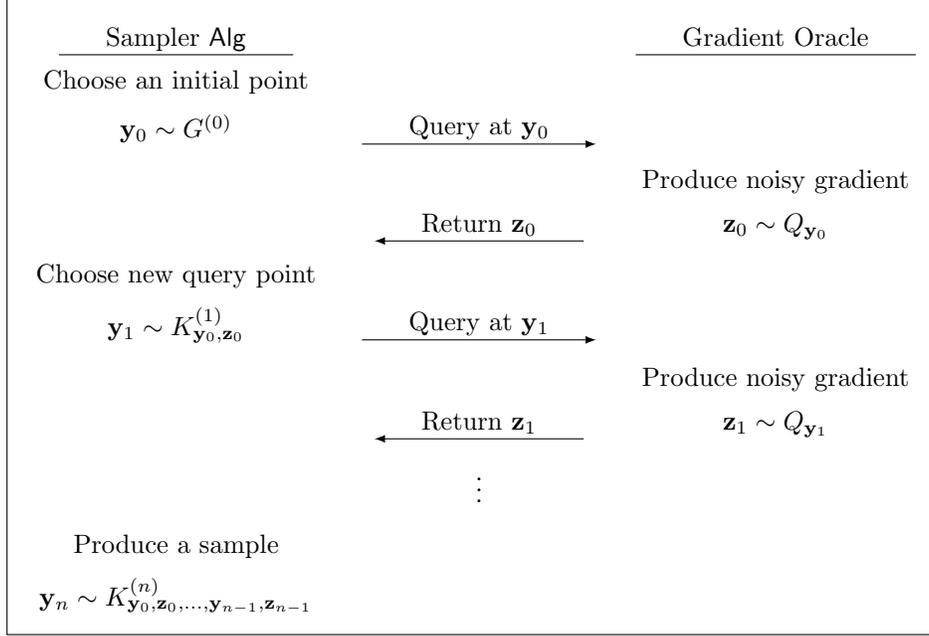
In this paper we will only be using measurable spaces that are product spaces of the reals. So from this point on we avoid explicitly mentioning the associated σ -field and always work with the appropriate Borel σ -field.

We use probability kernels to describe the response of the algorithm during its interaction with the gradient oracle.

Definition 3. A gradient-based sampling algorithm Alg on \mathbb{R}^d consists of the following elements:

1. A distribution $G^{(0)}$ over \mathbb{R}^d from which the initial point y_0 is drawn.
2. Probability kernels $K^{(i)}$ for $i \in \{1, \dots, n\}$, with the i^{th} probability kernel $K^{(i)}$ being a map from $\mathbb{R}^{d \times 2i}$ to probability measures over \mathbb{R}^d . The kernel $K^{(i)}$ is a map from the history of queries and gradients seen up to round i to a distribution over the next query point. The final kernel produces the sample.

To understand why this amounts to a valid and interesting class of sampling algorithms, let us understand the protocol to produce a sample given a stochastic gradient oracle Q (we use Q_x or $Q(\cdot|x)$ to denote the distribution of the gradients when queried at a point x).



Given the protocol above it is possible to write down the joint probability distribution of the path. For the sake of intuition, let us assume that the path of the algorithm has a density, and let $G^{(0)}$, Q , and $K^{(i)}$ also denote the corresponding densities¹. Then we can write the density of the path $\mathbf{y}_0 = y_0, \mathbf{z}_0 = z_0, \dots, \mathbf{y}_n = y_n$ as

$$G^{(0)}(y_0) \cdot Q(z_0|y_0) \cdot K^{(1)}(y_1|y_0, z_0) \cdot Q(z_1|y_1) \cdots K^{(n)}(y_n|y_{n-1}, z_{n-1}, \dots, y_0, z_0).$$

Let the distribution of the sample $\mathbf{y}_n \in \mathbb{R}^d$ produced by the interaction between the algorithm Alg and a gradient oracle Q be denoted by $\text{Alg}[n; Q]$. Explicitly, the marginal density of $\mathbf{y}_n = y_n$ is given by:

$$\begin{aligned} & \text{Alg}[n; Q](y_n) \\ & := \int G^{(0)}(y_0) Q(z_0|y_0) \cdots Q(z_{n-1}|y_{n-1}) K^{(n)}(y_n|y_{n-1}, z_{n-1}, \dots) dy_0 dz_0 \cdots dz_{n-1}. \end{aligned}$$

Restriction to finite parametric classes. As we are interested in demonstrating a lower bound for sampling algorithms, it suffices to work with only a finite indexed class of distributions $\{p_\theta\}_{\theta \in \Theta}$, where each p_θ is α -log smooth and $\alpha/2$ -strongly log-concave. Let f_θ be the negative log density corresponding to the distribution p_θ .

In the discussion that follows we let Q^θ be a stochastic gradient oracle that outputs $\nabla f_\theta(x) + \boldsymbol{\xi}$, where $\boldsymbol{\xi}$ is a Gaussian random variable with mean zero and covariance Σ_{noise} when queried at a point $x \in \mathbb{R}^d$. The covariance matrix Σ_{noise} will be specified

¹We note that our results hold when $G^{(0)}, K^{(1)}, \dots$ are arbitrary distributions.

subsequently. It suffices to lower bound the quantity on the right hand side below, which depends on this finite class of gradient oracles:

$$\inf_{\text{Alg}} \sup_Q \sup_{p^*} \text{TV}(\text{Alg}[n; Q], p^*) \geq \inf_{\text{Alg}} \sup_{\theta \in \Theta} \text{TV}(\text{Alg}[n; Q^\theta], p_\theta), \quad (2.1)$$

where the infimum over Alg is over the class of gradient-based sampling algorithms on \mathbb{R}^d .

Definition 4. An estimator μ on \mathbb{R}^d is a probability kernel from \mathbb{R}^d to the decision space \mathbb{R}^d .

An estimator μ , when given a sample x from the distribution p_θ , yields a distribution μ_x over the decision space \mathbb{R}^d . We choose the decision space to be \mathbb{R}^d which is distinct from the set of parameters Θ as we will be interested in estimates for the mean of p_θ .

Further, we define sequential randomized estimators—distributions over the decision space \mathbb{R}^d that arise by appropriately combining an estimator with a sampling algorithm.

Definition 5 (Sequential estimators). Given an estimator μ on \mathbb{R}^d , a gradient-based sampling algorithm Alg on \mathbb{R}^d , and a gradient oracle Q^θ , define a sequential estimator $\rho\{\mu, \text{Alg}[n; Q^\theta]\}$ as

$$\rho\{\mu, \text{Alg}[n; Q^\theta]\}(S) := \mathbb{E}_{\mathbf{y} \sim \text{Alg}[n; Q^\theta]} [\mu(S|\mathbf{y})], \quad \text{for all } S \in \mathcal{B}(\mathbb{R}^d).$$

Often we will simply use ρ to refer to the sequential estimator $\rho\{\mu, \text{Alg}[n; Q^\theta]\}$ that is built using the estimator μ and sampling algorithm Alg. In the remainder of the paper, when we write \inf_ρ , this indicates an infimum over both estimators μ and algorithms Alg that are used to build ρ .

A loss function $L = \{L_\theta : \theta \in \Theta\}$ on \mathbb{R}^d is a family of $[0, 1]$ -valued measurable functions on \mathbb{R}^d . Each L maps a $\theta \in \Theta$ into L_θ where $L_\theta(t)$ is the loss we suffer if we take the decision $t \in \mathbb{R}^d$ and θ is the true state of nature.

3. A characterization in terms of Bayes risk functions

In this section we prove a proposition that is crucial to prove our lower bound on the iteration complexity. Operationally, it reduces the problem of lower bounding the total variation distance between $\text{Alg}[n; Q^\theta]$ and p_θ to lower bounding a *difference of two Bayes risk functions*. This is advantageous as there exists a large set of well-developed techniques to both lower and upper bound Bayes risk functions.

As will be clear in the proof, the choice of total variation distance as our metric is pivotal. The total variation distance lends itself to a variational definition in terms of bounded loss functions, which is exploited in the proof of this proposition.

Proposition 1. (Randomization criterion) Let Θ be a finite set. Then for any loss function L_θ on \mathbb{R}^d we have

$$\begin{aligned} & \inf_{\text{Alg}} \sup_{\theta \in \Theta} \text{TV}(\text{Alg}[n; Q^\theta], p_\theta) \\ & \geq \inf_{\rho} \frac{1}{|\Theta|} \sum_{\theta \in \Theta} \mathbb{E}_{\mathbf{t} \sim \rho\{\mu, \text{Alg}[n; Q^\theta]\}} [L_\theta(\mathbf{t})] - \inf_{\mu} \frac{1}{|\Theta|} \sum_{\theta \in \Theta} \mathbb{E}_{\mathbf{x} \sim p_\theta} [\mathbb{E}_{\mathbf{t} \sim \mu_{\mathbf{x}}} [L_\theta(\mathbf{t})]], \end{aligned}$$

where the infimum over Alg is over gradient-based sampling algorithms, the infimum over ρ is as above, and the infimum over μ is over estimators on \mathbb{R}^d .

As mentioned earlier, the infimum over sequential estimator ρ in the proposition above involves taking an infimum simultaneously over both estimators μ and sampling algorithms Alg that make up ρ . The proof of the proposition is in Appendix A.

To understand why this proposition is useful let us unpack the right hand side of the randomization criterion which is a difference of two Bayes risk functions. Given a fixed loss function L_θ , the first term is the average risk associated with the optimal estimator that is fed a *fictitious* sample generated by the optimal sampling algorithm after n rounds of interaction with the stochastic gradient oracle. This is compared to the second term which is the average risk of an optimal estimator that receives a single *real* sample from the target distribution p_θ .

Thus, to demonstrate a lower bound for sampling we will construct a pair of two distributions p_1 and p_2 , and a gradient oracle such that:

1. It is extremely easy to estimate the mean of the distribution given a single sample from the real distribution. This is to ensure that the second term is small and is independent of dimension.
2. But, the Bayes risk associated with estimating the mean when given a fictitious sample from the sampler grows at a rate of $\Omega(\sigma\sqrt{d/n})$.

Such a construction ensures that the difference between these Bayes risk functions is large and proves a lower bound on the iteration complexity of sampling via the randomization criterion.

The term on the left hand side of our randomization criterion is similar in spirit to a quantity called the Le Cam deficiency between two families of distributions. Given two families $\{R_\theta\}_{\theta \in \Theta}$ and $\{S_\theta\}_{\theta \in \Theta}$, the Le Cam deficiency between these families is defined to be

$$\inf_M \sup_{\theta \in \Theta} \text{TV}(MR_\theta, S_\theta),$$

where M is a probability kernel which takes as input samples from R_θ . It is measure of how difficult it is to transform a sample from the distribution R_θ to a sample from S_θ using a probability kernel M (which does not vary with θ). In our setting, in some sense we want to transform n adaptively chosen samples from the gradient oracle Q^θ to a sample from the target distribution p_θ .

The proposition above is similar in flavour to the classical Le Cam randomization criterion [24] studied extensively in the literature on the comparison of statistical experiments [see, e.g, 47, 25]. In fact, it is a version of Le Cam’s randomization criterion restricted to a special class of sequential estimators.

Finally, we note that an object similar to the first Bayes risk function (when the estimator is fed fictitious samples from the sampling algorithm) also arises in the study of lower bounds on the iteration complexity of gradient-based stochastic optimization algorithms. At a high level, there, this Bayes risk lower bounds the worst case gap to the optimum for the optimal stochastic optimization algorithm. We point the interested reader to the work of [39], [1], and [40] for further details. Interestingly, in the study of lower bounds for stochastic optimization algorithms, no term analogous to the second Bayes risk function (when the estimator is fed real samples from p_θ) arises, which adds a level of difficulty to our problem.

4. Lower bounds for stochastic first order methods

We now state our main result for stochastic gradient oracles.

Theorem 4.1. *For all $d, \sigma^2, n \geq \sigma^2 d/4$ and for all $\alpha \leq \sigma^2 d/(256n)$,*

$$\inf_{\text{Alg}} \sup_Q \sup_{p^*} \text{TV}(\text{Alg}[n; Q], p^*) \geq \frac{\sigma}{16} \sqrt{\frac{d}{n}},$$

where the infimum is over all gradient-based sampling algorithms, the supremum over Q is over unbiased gradient oracles with noise-variance bounded by $\sigma^2 d$, and the supremum over p^* is over α -log smooth, $\alpha/2$ -strongly log-concave distributions over \mathbb{R}^d with the mean $\|\mathbb{E}_{\mathbf{x} \sim p^*} [\mathbf{x}]\|_2 \leq \frac{1}{\alpha}$.

The interpretation of the theorem given the order of the quantifiers, $\inf_{\text{Alg}} \sup_Q \sup_{p^*}$ is as follows: Given any d, σ^2 and n , for every gradient-based sampling algorithm, there exists a noise oracle and a target distribution (both of which could depend on the sampling algorithm and on the problem parameters d, σ^2 and n) such that the total variation between the distribution of points generated by the algorithm and the target distribution is lower bounded by $c\sigma\sqrt{d/n}$.

To reiterate, we only consider distributions that have their means in a ball of radius $1/\alpha$ around the origin. This ensures that most of the mass of these distributions is concentrated in a ball of radius $\mathcal{O}(\sqrt{d/\alpha})$ around the mean (as p^* is $\alpha/2$ -strongly log-concave and hence sub-Gaussian). This constraint on the family of target distributions ensures that the lower bound does not trivially hold by choosing an appropriate family of distributions with means that are arbitrarily far away from each other.

We note that our lower bound only holds for the case when the gradients are noisy. If we are given exact gradients there exist methods that use a Metropolis-Hastings filter that provably converge exponentially fast (in the target accuracy) when the target distribution is strongly log-concave [see, e.g., 15].

It is also fairly straightforward to run through our argument and also arrive at a lower bound in terms of Kullback-Leibler divergence by using the Pinsker-Csiszár-Kullback inequality.

Proof of Corollary 1 and its consequences. If we apply Theorem 4.1 for $n = \sigma^2 d/4$, then the claim of the corollary follows immediately by choosing $\alpha = 1/64$.

Corollary 1 dictates that there exists a family of well-conditioned distributions with *constant* log-smoothness such that it takes at least order $\sigma^2 d$ queries to the stochastic gradient oracle to obtain a sample whose distribution is a constant $(1/8)$ away from the target distribution in total variation distance.

When do we have optimal algorithms? One candidate would be the Stochastic Gradient Langevin Dynamics (SGLD) algorithm [49], which is a version of the ULA that uses stochastic gradients instead of exact ones. Corollary 25 in [13] characterizes an upper bound for the rate of convergence of SGLD in total variation distance. Let $\text{SGLD}[n; Q]$ denote the distribution of the sample produced by SGLD (with parameters tuned according to requirements of their results) after n iterations. Then if $\alpha, \sigma^2 \asymp c$ and the distribution is well-conditioned their results guarantee,

$$\text{TV}(\text{SGLD}[n; Q], p^*) \leq \tilde{\mathcal{O}} \left(\sqrt{\frac{d}{n}} \right).$$

If σ^2 and α are viewed as constants, this matches our lower bound on the number of iterations needed to achieve constant accuracy.

Proof details for Theorem 4.1

Before we start our proof of the main theorem, we set up some quantities that will be useful in the proof. Let $\lambda \in [0, 1/2]$ be a scalar parameter that will be specified in the sequel. We consider a family of two Gaussian distributions,

$$p_\theta := \mathcal{N}(m_\theta, I_{d \times d}/\alpha), \quad \text{where } \theta \in \Theta = \{1, 2\}. \quad (4.1)$$

The last $d - 1$ coordinates of the two mean vectors are all 0. The mean vectors differ only on the first coordinate. Set $m_{11} = \lambda/\alpha$ and $m_{21} = -\lambda/\alpha$, where m_{11} and m_{21} refers to the first coordinate of the vectors. Clearly, these vectors are separated in Euclidean norm

$$\|m_1 - m_2\|_2 = \frac{2\lambda}{\alpha}. \quad (4.2)$$

The distributions we have defined are Gaussian, therefore, each distribution in this family is α -log smooth and $\alpha/2$ -strongly log-concave (in fact, it is α -strongly log-concave).

We choose our decision space to be the space of vectors in \mathbb{R}^d and define a loss function over this space as,

$$L_\theta(t) := \min \{ \alpha \|t - m_\theta\|_2, 1 \}, \quad \text{for all } \theta \in \Theta. \quad (4.3)$$

We will be working with a fixed gradient oracle Q^θ that adds zero mean Gaussian noise with covariance

$$\Sigma_{noise} = \begin{bmatrix} \sigma^2 d & 0_{d-1} \\ 0_{d-1}^\top & 0_{(d-1) \times (d-1)} \end{bmatrix},$$

to the gradient of f_θ when queried at x . Since $\text{Tr}(\Sigma_{noise}) = \sigma^2 d$, this is a valid gradient oracle.

This noise oracle uses its entire noise budget on the first coordinate and adds Gaussian noise with variance $\sigma^2 d$ to the gradient along this coordinate. We chose such an oracle because the two distributions p_1 and p_2 differ only along this single coordinate. [The choice of both the distributions and the noise oracle is guided by Proposition 1.](#) We will establish our lower bound by lower bounding a difference of Bayes risk functions. The first Bayes risk function shall correspond to a problem where the learner needs to estimate the first coordinate of the mean of the distribution when given access to n stochastic gradient queries of the log-density. Since the noise oracle packs all of the variance into the first coordinate, this problem is challenging and the first Bayes risk function will be relatively large. The second Bayes risk function corresponds to a problem where the learner is asked to estimate the first coordinate of the mean using just a single sample from the true distribution. As the means of p_1 and p_2 differ only along the first coordinate, the estimation of the mean using just a single sample from the underlying distribution is feasible, which means that the value of the second Bayes risk function will be relatively small. This shall guarantee that the difference between the two Bayes risk functions is large. We are now ready to begin the proof.

Proof. : We want a lower bound on the quantity,

$$\inf_{\text{Alg}} \sup_Q \sup_{p^*} \text{TV}(\text{Alg}[n; Q], p^*).$$

First we replace the supremum over all distributions p^* with our pair of Gaussian distributions parameterized by Θ defined in Equation (4.1) above. Further, as mentioned above we drop the supremum over gradient oracle distributions and consider the gradient oracle Q^θ that adds Gaussian noise to the first coordinate of the gradient. By relaxing the problem in this manner we must have,

$$\mathcal{M}_n := \inf_{\text{Alg}} \sup_Q \sup_{p^*} \text{TV}(\text{Alg}[n; Q], p^*) \geq \inf_{\text{Alg}} \sup_{\theta \in \Theta} \text{TV}(\text{Alg}[n; Q^\theta], p_\theta).$$

Let μ denote an estimator (see Definition 4) and ρ denote a sequential estimator (see

Definition 5). Then by our randomization criterion (Proposition 1),

$$\begin{aligned} \mathcal{M}_n &\geq \inf_{\text{Alg}} \sup_{\theta \in \Theta} \text{TV}(\text{Alg}[n; Q^\theta], p_\theta) \\ &\geq \underbrace{\inf_{\rho} \frac{1}{|\Theta|} \sum_{\theta \in \Theta} \mathbb{E}_{\mathbf{t} \sim \rho\{\mu, \text{Alg}[n; Q^\theta]\}} [L_\theta(\mathbf{t})]}_{=:\Psi_1} - \underbrace{\inf_{\mu} \frac{1}{|\Theta|} \sum_{\theta \in \Theta} \mathbb{E}_{\mathbf{x} \sim p_\theta} [\mathbb{E}_{\mathbf{t} \sim \mu_{\mathbf{x}}} [L_\theta(\mathbf{t})]]}_{=:\Psi_2}. \end{aligned} \quad (4.4)$$

To lower bound \mathcal{M}_n we need to establish a lower bound on Ψ_1 and an upper bound on Ψ_2 .

The term Ψ_1 is the Bayes risk of the optimal sequential estimator ρ with respect to the loss function L_θ when the true parameter θ is drawn from a uniform distribution over the set $\Theta = \{1, 2\}$. To prove a lower bound on this term we turn to Le Cam's method. Roughly, this entails reducing this mean estimation problem to a hypothesis testing problem via a standard argument (see Lemma 1). Then, we lower bound the error made in this hypothesis testing problem by the optimal test. This is achieved by upper bounding the total variation distance between the distributions of the algorithm's history over n rounds when the true distribution is either p_1 or p_2 (in Lemma 2). This argument leads to the lower bound: $\Psi_1 > \lambda(1 - \lambda\sqrt{n}/(\sigma\sqrt{d}))$. The details for this lower bound are presented in Appendix B.1.

Similar to the first term, Ψ_2 is the Bayes risk of the optimal estimator μ that is given a *single* sample from the true distribution p_θ . A bound of $\Psi_2 < \sqrt{\alpha}$ follows by considering the estimator that outputs the first coordinate of the sample along the first dimension (the standard deviation along with dimension is $1/\sqrt{\alpha}$ and $L_\theta(t) = \alpha\|t - m_\theta\|_2$) and 0 along the remaining dimensions. This simple calculation is detailed in Appendix B.2.

By combining these results,

$$\mathcal{M}_n \geq \lambda \left(1 - \frac{\lambda}{\sigma} \sqrt{\frac{n}{d}} \right) - \sqrt{\alpha}.$$

Pick $\lambda = \sigma\sqrt{d}/(4\sqrt{n})$. Recall that the number of queries $n \geq \sigma^2 d/4$, which guarantees that $\lambda \leq 1/2$. Plugging this value of λ into the inequality above,

$$\mathcal{M}_n \geq \frac{3\sigma}{16} \sqrt{\frac{d}{n}} - \sqrt{\alpha} \geq \frac{\sigma}{16} \sqrt{\frac{d}{n}},$$

where the condition on the smoothness $\alpha \leq \sigma^2 d/(256n)$ ensures that the second term in the inequality above is smaller than the first, which completes the proof. \square

5. Discussion and open problems

Many interesting questions remain open related to lower bounding the iteration complexity of sampling algorithms. In our lower bounds we ignored the dependence on the

condition number $\kappa = \alpha/\beta$, an important parameter of the problem, and established a lower bound for well-conditioned distributions (where $\kappa = 2$). It would be interesting to study if the lower bounds can be extended to also capture the dependence on κ . Another related question concerns the iteration complexity of sampling when the potentials are smooth and strongly-convex with respect to general ℓ_p norms for $p \in [1, \infty)$.

We have given a lower bound in terms of the variance, number of variables, and number of samples, for a worst-case smooth distribution. The degree of smoothness in our construction depends (albeit somewhat mildly) on the other parameters. It would be interesting to study how the optimal accuracy behaves when the smoothness and the other parameters are varied independently.

In our construction, all of the noise in the gradients is concentrated in one component. This noise satisfies the constraints used to prove upper bounds in papers such as [12, 10], so that our analysis rules out certain kinds of improvements to those bounds. On the other hand, the noise arising in applications, for example when stochastic gradients are obtained by subsampling, is often qualitatively unlike this. This raises the question of whether or not more efficient sampling is possible when noise satisfies constraints other than bounded variance, for example when it is isotropic, or nearly so.

After this work appeared in preliminary form [8], Johndrow et al. [21] established lower bounds for a large class of models and subsampled MCMC methods. It would be interesting to see if these results can be extended to obtain information-theoretic lower bounds similar to the ones in this paper.

Another direction is to develop techniques to establish lower bounds for sampling methods when we have access to the exact gradients of the potential function. On this front, there has been some recent progress on related problems, which may provide some hints as to how one might proceed. Ge et al. [19] provide a lower bound on the number of exact gradient queries required to estimate the normalization constant of strongly log-concave distributions. Rademacher et al. [38] showed that a quadratic number of queries to a deterministic membership oracle is required to estimate the volume of a convex body in d dimensions. Talwar [46] recently exhibited a family of distributions with non-convex potentials for which sampling is NP-Hard.

Appendix A: Proof of Proposition 1

Given any sampling algorithm Alg, for any estimator μ it is possible to define a sequential estimator $\rho\{\mu, \text{Alg}[n; Q^\theta]\}$ by using both Alg and μ . Then, for all $\theta \in \Theta$ we have,

$$\begin{aligned} & \frac{1}{|\Theta|} \sum_{\theta \in \Theta} (\mathbb{E}_{\mathbf{t} \sim \rho\{\mu, \text{Alg}[n; Q^\theta]\}} [L_\theta(\mathbf{t})] - \mathbb{E}_{\mathbf{x} \sim p_\theta} [\mathbb{E}_{\mathbf{t} \sim \mu_{\mathbf{x}}} [L_\theta(\mathbf{t})]]) \\ &= \frac{1}{|\Theta|} \sum_{\theta \in \Theta} (\mathbb{E}_{\mathbf{x} \sim \text{Alg}[n; Q^\theta]} [\mathbb{E}_{\mathbf{t} \sim \mu_{\mathbf{x}}} [L_\theta(\mathbf{t})]] - \mathbb{E}_{\mathbf{x} \sim p_\theta} [\mathbb{E}_{\mathbf{t} \sim \mu_{\mathbf{x}}} [L_\theta(\mathbf{t})]]) \end{aligned}$$

where the equality above follows by the definition of the sequential estimator ρ . Define a function $h_\theta(x) := \mathbb{E}_{\mathbf{t} \sim \mu_x} [L_\theta(\mathbf{t})]$. Clearly this function is also bounded, $0 \leq h_\theta \leq 1$ since

L_θ is bounded between 0 and 1. Hence, we have,

$$\begin{aligned}
& \frac{1}{|\Theta|} \sum_{\theta \in \Theta} (\mathbb{E}_{\mathbf{t} \sim \rho\{\mu, \text{Alg}[n; Q^\theta]\}} [L_\theta(\mathbf{t})] - \mathbb{E}_{\mathbf{x} \sim p_\theta} [\mathbb{E}_{\mathbf{t} \sim \mu_{\mathbf{x}}} [L_\theta(\mathbf{t})]]) \\
&= \frac{1}{|\Theta|} \sum_{\theta \in \Theta} (\mathbb{E}_{\mathbf{x} \sim \text{Alg}[n; Q^\theta]} [h_\theta(\mathbf{x})] - \mathbb{E}_{\mathbf{x} \sim p_\theta} [h_\theta(\mathbf{x})]) \\
&\leq \frac{1}{|\Theta|} \sum_{\theta \in \Theta} \sup_{h_\theta: 0 \leq h_\theta \leq 1} (\mathbb{E}_{\mathbf{x} \sim \text{Alg}[n; Q^\theta]} [h_\theta(\mathbf{x})] - \mathbb{E}_{\mathbf{x} \sim p_\theta} [h_\theta(\mathbf{x})]) \\
&\stackrel{(i)}{=} \frac{1}{|\Theta|} \sum_{\theta \in \Theta} \text{TV}(\text{Alg}[n; Q^\theta], p_\theta) \\
&\leq \sup_{\theta \in \Theta} \text{TV}(\text{Alg}[n; Q^\theta], p_\theta),
\end{aligned}$$

where the equality (i) is due to the variational definition of the total variation distance. Taking an infimum over all possible sequential estimators ρ (by taking an infimum over all estimators ν and algorithms $\widetilde{\text{Alg}}$) we find that for all estimators μ ,

$$\begin{aligned}
& \sup_{\theta \in \Theta} \text{TV}(\text{Alg}[n; Q^\theta], p_\theta) \\
&\geq \inf_{\nu, \widetilde{\text{Alg}}} \frac{1}{|\Theta|} \sum_{\theta \in \Theta} (\mathbb{E}_{\mathbf{t} \sim \rho\{\nu, \widetilde{\text{Alg}}[n; Q^\theta]\}} [L_\theta(\mathbf{t})] - \mathbb{E}_{\mathbf{x} \sim p_\theta} [\mathbb{E}_{\mathbf{t} \sim \mu_{\mathbf{x}}} [L_\theta(\mathbf{t})]]) \\
&= \inf_{\nu, \widetilde{\text{Alg}}} \left[\frac{1}{|\Theta|} \sum_{\theta \in \Theta} \mathbb{E}_{\mathbf{t} \sim \rho\{\nu, \widetilde{\text{Alg}}[n; Q^\theta]\}} [L_\theta(\mathbf{t})] \right] - \frac{1}{|\Theta|} \sum_{\theta \in \Theta} \mathbb{E}_{\mathbf{x} \sim p_\theta} [\mathbb{E}_{\mathbf{t} \sim \mu_{\mathbf{x}}} [L_\theta(\mathbf{t})]].
\end{aligned}$$

As the above inequality holds for all estimators μ and algorithms Alg , taking a supremum over μ on the right side and an infimum over all gradient-based sampling algorithms on the left hand side completes the proof.

Appendix B: Additional proof details for lower bounds with gradient oracles

B.1. Lower Bound on Ψ_1

In this subsection we prove a lower bound on Ψ_1 . By definition

$$\Psi_1 = \inf_{\rho} \left\{ \frac{1}{|\Theta|} \sum_{\theta \in \Theta} \mathbb{E}_{\mathbf{t} \sim \rho\{\mu, \text{Alg}[n; Q^\theta]\}} [L_\theta(\mathbf{t})] \right\}.$$

This is the Bayes risk of sequential estimators with respect to the loss function L_θ defined in Equation (4.3). We want a lower bound on it, which is a lower bound on the risk of

the optimal estimator that attempts to identify the underlying value of the mean vector associated with θ given n adaptive queries to a stochastic gradient oracle.

Let random variables in a sample path of the algorithm be $\mathbf{y}_0, \mathbf{z}_0, \dots, \mathbf{y}_{n-1}, \mathbf{z}_{n-1}, \mathbf{y}_n, \mathbf{t}$, where $(\mathbf{y}_1, \dots, \mathbf{y}_{n-1})$ are the query points of the algorithm, $(\mathbf{z}_0, \dots, \mathbf{z}_{n-1})$ are the stochastic gradients returned by the oracle, \mathbf{y}_n is the sample produced and \mathbf{t} is the estimate of the mean. Define

$$\mathbf{\Pi} := \{\mathbf{y}_0, \mathbf{z}_0, \mathbf{y}_1, \mathbf{z}_1, \dots, \mathbf{y}_{n-1}, \mathbf{z}_{n-1}, \mathbf{y}_n, \mathbf{t}\}.$$

In Lemma 1 below, we prove that

$$\Psi_1 \geq \lambda (1 - \text{TV}(\mathbb{P}_n^1, \mathbb{P}_n^2)),$$

where \mathbb{P}_n^1 is the distribution of $\mathbf{\Pi}$ when $\theta = 1$ and \mathbb{P}_n^2 is its distribution when $\theta = 2$. The proof uses a fairly standard argument (called Le Cam's method) to reduce from mean estimation to testing [see, e.g., 48, Chapter 15]. To complete our lower bound, we need to upper bound the total variation distance between the distributions \mathbb{P}_n^1 and \mathbb{P}_n^2 . We do this by instead controlling the Kullback-Leibler divergence between these distributions, which bounds the total variation distance. See Lemma 2 below for the details of this calculation. This upper bound combined with the inequality in the display above yields

$$\Psi_1 \geq \lambda \left(1 - \frac{\lambda}{\sigma} \sqrt{\frac{n}{d}}\right), \quad (\text{B.1})$$

which is the desired lower bound on Ψ_1 .

Auxillary lemmas

Lemma 1. *Let the Bayes risk Ψ_1 be as defined in (4.4), and \mathbb{P}_n^1 and \mathbb{P}_n^2 be as defined above, then,*

$$\Psi_1 \geq \lambda (1 - \text{TV}(\mathbb{P}_n^1, \mathbb{P}_n^2)).$$

Proof. By the definition of a sequential estimator ρ in Definition 5,

$$\rho\{\mu, \text{Alg}[n; Q^\theta]\}(\cdot) := \mathbb{E}_{\mathbf{y} \sim \text{Alg}[n; Q^\theta]} [\mu(\cdot | \mathbf{y})].$$

It is made up of an estimator μ and a sampling algorithm Alg . Let $\mathbf{\Pi}$ be a sample path of the algorithm, as defined above.

Reduction to testing: We proceed by the standard reduction used in proving lower bounds on Bayes risk functions, that of reducing from an estimation problem to a testing problem [see, e.g., 48, Chapter 15]. Recall the definition of Ψ_1

$$\Psi_1 = \inf_{\rho} \left\{ \frac{1}{|\Theta|} \sum_{\theta \in \Theta} \mathbb{E}_{\mathbf{t} \sim \rho\{\mu, \text{Alg}[n; Q^\theta]\}} [L_\theta(\mathbf{t})] \right\}.$$

For a fixed ρ , if $\|\mathbf{t} - m_\theta\|_2 \geq \lambda/\alpha$ then

$$L_\theta(\mathbf{t}) = \min\{\alpha\|\mathbf{t} - m_\theta\|_2, 1\} \geq \min\{\lambda, 1\} = \lambda,$$

as $\lambda \in [0, 1/2]$. Thus, by applying Markov's inequality with respect to the distribution ρ we infer,

$$\begin{aligned} \Psi_1 &= \inf_{\rho} \left\{ \frac{1}{|\Theta|} \sum_{\theta \in \Theta} \mathbb{E}_{\mathbf{t} \sim \rho\{\mu, \text{Alg}[n; Q^\theta]\}} [L_\theta(\mathbf{t})] \right\}, \\ &\geq \lambda \cdot \inf_{\rho} \underbrace{\left\{ \frac{1}{|\Theta|} \sum_{\theta \in \Theta} \mathbb{E}_{\mathbf{t} \sim \rho\{\mu, \text{Alg}[n; Q^\theta]\}} \left[\mathbb{I} \left[\|\mathbf{t} - m_\theta\| \geq \frac{\lambda}{\alpha} \right] \right] \right\}}_{=: \Xi(\rho)}, \end{aligned} \quad (\text{B.2})$$

where $\Xi(\rho)$ is the expectation of the indicator function under the distribution ρ over t and under the uniform mixture distribution over θ . Let $\mathbb{P}_{\text{joint}}$ denote this joint probability over the random variables $\mathbf{\Pi}$ and \mathbf{j} (which takes values in Θ) then,

$$\frac{1}{|\Theta|} \sum_{\theta \in \Theta} \mathbb{E}_{\mathbf{t} \sim \rho\{\mu, \text{Alg}[n; Q^\theta]\}} \left[\mathbb{I} \left[\|\mathbf{t} - m_\theta\| \geq \frac{\lambda}{\alpha} \right] \right] =: \mathbb{P}_{\text{joint}} \left(\|\mathbf{t} - m_{\mathbf{j}}\|_2 \geq \frac{\lambda}{\alpha} \right).$$

So we have reduced it to lower bounding the quantity on the right hand side. Toward this end, let us consider a test ϕ

$$\phi(\mathbf{\Pi}) := \underset{j \in \Theta}{\text{argmin}} \{ \|\mathbf{t} - m_j\|_2 \},$$

where we break ties arbitrarily. Conditioned on $\mathbf{j} = \theta$: we claim that the event $\|\mathbf{t} - m_\theta\|_2 < \lambda/\alpha$ ensures that the test $\phi(\mathbf{\Pi})$ is correct and returns θ . In order to see this, note that for any other $\theta' \in \Theta$, by triangle inequality we have,

$$\|m_{\theta'} - \mathbf{t}\|_2 \geq \|m_\theta - m_{\theta'}\|_2 - \|m_\theta - \mathbf{t}\|_2 > \frac{2\lambda}{\alpha} - \frac{\lambda}{\alpha} = \frac{\lambda}{\alpha}.$$

The lower bound $\|m_\theta - m_{\theta'}\|_2 \geq 2\lambda/\alpha$ follows by our construction of the set of mean vectors, see Inequality (4.2). As these vectors are separated we must also have, $\|m_{\theta'} - \mathbf{t}\|_2 \geq \|m_\theta - \mathbf{t}\|_2$, where $\theta' \neq \theta$.

Therefore, conditioned on a value of the random variable $\mathbf{j} = \theta$, the event $\{\|\mathbf{t} - m_\theta\|_2 < \lambda/\alpha\}$ is contained within the event $\{\phi(\mathbf{\Pi}) = \theta\}$. This implies,

$$\mathbb{E}_{\mathbf{t} \sim \rho\{\mu, \text{Alg}[n; Q^\theta]\}} \left[\mathbb{I} \left[\|\mathbf{t} - m_\theta\| \geq \frac{\lambda}{\alpha} \right] \right] \geq \mathbb{E}_{(\mathbf{\Pi}, \mathbf{j}) \sim \mathbb{P}_{\text{joint}}} [\mathbb{I} [\phi(\mathbf{\Pi}) \neq \mathbf{j}]].$$

Taking expectations with respect to the uniform mixture distribution over θ yields,

$$\Xi(\rho) \geq \mathbb{P}_{\text{joint}}(\phi(\mathbf{\Pi}) \neq \mathbf{j}).$$

We take an infimum over all sequential estimators ρ and an infimum over all tests ϕ , coupled with Inequality (B.2) to conclude that

$$\Psi_1 \geq \lambda \cdot \inf_{\rho} \Xi(\rho) \geq \lambda \inf_{\phi} \mathbb{P}_{\text{joint}}(\phi(\mathbf{\Pi}) \neq \mathbf{j}). \quad (\text{B.3})$$

Le Cam's Method: The next step is to demonstrate a lower bound on failure probability of any test ϕ for this, we use Le Cam's method [see, e.g., 48, Section 15.2.1],

$$\inf_{\phi} \mathbb{P}_{\text{joint}}(\phi(\mathbf{\Pi}) \neq \mathbf{j}) = 1 - \text{TV}(\mathbb{P}_n^1, \mathbb{P}_n^2),$$

where \mathbb{P}_n^1 and \mathbb{P}_n^2 denote the distribution of $\mathbf{\Pi}$ when the true underlying target distribution is p_1 and p_2 respectively. Combining this lower bound with Inequality (B.3) completes the proof. \square

The next lemma establishes an upper bound on the total variation distance between the distributions \mathbb{P}_n^1 and \mathbb{P}_n^2 .

Lemma 2. *Let $\mathbf{\Pi}$, \mathbb{P}_n^1 and \mathbb{P}_n^2 be as defined above. Then*

$$\text{TV}(\mathbb{P}_n^1, \mathbb{P}_n^2) \leq \frac{\lambda}{\sigma} \sqrt{\frac{n}{d}}.$$

Proof. Recall that the random variable $\mathbf{\Pi} = \{\mathbf{y}_0, \mathbf{z}_0, \mathbf{y}_1, \mathbf{z}_1, \dots, \mathbf{y}_{n-1}, \mathbf{z}_{n-1}, \mathbf{y}_n, \mathbf{t}\}$ is comprised of the query points of the algorithm $(\mathbf{y}_1, \dots, \mathbf{y}_{n-1})$, the stochastic gradients $(\mathbf{z}_0, \dots, \mathbf{z}_{n-1})$, the output of the sampling algorithm \mathbf{y}_n and the estimate for the mean vector \mathbf{t} . The distribution \mathbb{P}_n^1 is the law of $\mathbf{\Pi}$ when $\theta = 1$ and \mathbb{P}_n^2 is its law when $\theta = 2$.

Recall that the oracle noise distribution $Q_y^\theta \sim \mathcal{N}(\nabla f_\theta(y), \Sigma_{\text{noise}})$, adds mean-zero Gaussian noise to the gradient of f_θ where

$$\Sigma_{\text{noise}} = \begin{bmatrix} \sigma^2 d & 0_{d-1} \\ 0_{d-1}^\top & 0_{(d-1) \times (d-1)} \end{bmatrix}.$$

By this choice of the noise oracle the probability measure \mathbb{P}_n^1 is absolutely continuous² with respect to \mathbb{P}_n^2 . This allows us to bound the Kullback-Leibler divergence between \mathbb{P}_n^1 and \mathbb{P}_n^2 , which also leads to a bound on the total variation distance between them.

$$\begin{aligned} & \text{KL}(\mathbb{P}_n^1 | \mathbb{P}_n^2) \\ &= \mathbb{E}_{\mathbf{\Pi} \sim \mathbb{P}_n^1} \left[\log \left(\frac{G^{(0)}(\mathbf{y}_0) Q^1(\mathbf{z}_0 | \mathbf{y}_0) K^{(1)}(\mathbf{y}_1 | \mathbf{y}_0, \mathbf{z}_0) \cdots K^{(n)}(\mathbf{y}_n | \mathbf{y}_{n-1}, \dots) \mu(\mathbf{t} | \mathbf{y}_n)}{G^{(0)}(\mathbf{y}_0) Q^2(\mathbf{z}_0 | \mathbf{y}_0) K^{(1)}(\mathbf{y}_1 | \mathbf{y}_0, \mathbf{z}_0) \cdots K^{(n)}(\mathbf{y}_n | \mathbf{y}_{n-1}, \dots) \mu(\mathbf{t} | \mathbf{y}_n)} \right) \right] \\ &= \mathbb{E}_{\mathbf{\Pi} \sim \mathbb{P}_n^1} \left[\log \left(\frac{\prod_{t=0}^{n-1} Q^1(\mathbf{z}_t | \mathbf{y}_t)}{\prod_{t=0}^{n-1} Q^2(\mathbf{z}_t | \mathbf{y}_t)} \right) \right] \\ &= \sum_{t=0}^{n-1} \mathbb{E}_{\mathbf{y}_0, \mathbf{z}_0, \dots, \mathbf{y}_t \sim \mathbb{P}_n^1} \left[\mathbb{E}_{\mathbf{z}_t \sim Q_y^1} \left[\log \left(\frac{Q^1(\mathbf{z}_t | \mathbf{y}_t)}{Q^2(\mathbf{z}_t | \mathbf{y}_t)} \right) \middle| \mathbf{y}_0 = \mathbf{y}_0, \mathbf{z}_0 = \mathbf{z}_0, \dots, \mathbf{z}_{t-1} = \mathbf{z}_{t-1}, \mathbf{y}_t = \mathbf{y}_t \right] \right] \\ &\stackrel{(i)}{\leq} n \cdot \max_{y \in \mathbb{R}^d} \left\{ \mathbb{E}_{\mathbf{z} \sim Q_y^1} \left[\log \left(\frac{Q^1(\mathbf{z} | y)}{Q^2(\mathbf{z} | y)} \right) \right] \right\} \\ &\stackrel{(ii)}{=} \frac{n}{2} \max_{y \in \mathbb{R}^d} [\nabla f_1(y) - \nabla f_2(y)]^\top \Sigma_{\text{noise}}^\dagger [\nabla f_1(y) - \nabla f_2(y)], \end{aligned} \tag{B.4}$$

²See, for example, [22], for a review of absolute continuity.

where (i) follows as we upper bound the sum of the divergences of the conditional distributions by n times the divergence between the conditional distributions at the maximal point. The equality in (ii) follows by the formula for KL divergence between two Gaussians with the same covariance Σ_{noise} and different means (Σ_{noise}^\dagger is the Moore-Penrose inverse of Σ_{noise}). Given our definition of the target distributions p_θ (defined in Equation (4.1)), it can be verified that the negative logarithm of the density is

$$f_\theta(y) = \frac{\alpha}{2}(y - m_\theta)^\top (y - m_\theta), \quad \text{for } \theta \in \{1, 2\}.$$

Therefore, for any $y \in \mathbb{R}^d$,

$$\begin{aligned} & [\nabla f_1(y) - \nabla f_2(y)]^\top \Sigma_{noise}^\dagger [\nabla f_1(y) - \nabla f_2(y)] \\ &= [\alpha(y - m_1) - \alpha(y - m_2)]^\top \Sigma_{noise}^\dagger [\alpha(y - m_1) - \alpha(y - m_2)] \\ &= \alpha^2 [m_2 - m_1]^\top \Sigma_{noise}^\dagger [m_2 - m_1] = \frac{4\lambda^2}{\sigma^2 d}, \end{aligned}$$

the final equality is because $\|m_2 - m_1\|_2^2 = 4\lambda^2/\alpha^2$. This equation along with Inequality (B.4) yields

$$\text{KL}(\mathbb{P}_n^1 | \mathbb{P}_n^2) \leq \frac{2n\lambda^2}{\sigma^2 d}.$$

Finally, the Pinsker-Csiszár-Kullback inequality implies

$$\text{TV}(\mathbb{P}_n^1, \mathbb{P}_n^2) \leq \sqrt{\text{KL}(\mathbb{P}_n^1 | \mathbb{P}_n^2)/2} \leq \lambda\sqrt{n}/(\sigma\sqrt{d}),$$

which completes the proof. \square

B.2. Upper bound on Ψ_2

In this section we construct a simple estimator that is given access to a single sample from the true distribution p_θ , and we upper bound its risk. This immediately bounds the Bayes risk Ψ_2 .

Lemma 3. *The Bayes risk Ψ_2 defined in (4.4) satisfies*

$$\Psi_2 \leq \sqrt{\alpha}.$$

Proof. Recall the definition of the Gaussian distribution p_θ ,

$$p_\theta = \mathcal{N}(m_\theta, I_{d \times d}/\alpha).$$

Let \mathbf{x} be a sample from the Gaussian distribution p_θ which has mean m_θ and covariance $I_{d \times d}/\alpha$. The mean m_θ has all components except the first to be equal to 0. Therefore let

\hat{m} be an estimator that outputs the first coordinate of the sample it receives in the first coordinate and 0 in all other coordinates. Then we have,

$$\mathbb{E}_{\mathbf{x}} [\|\hat{m}(\mathbf{x}) - m_{\theta}\|_2] = \mathbb{E}_{\mathbf{x}} [|\hat{m}_1(\mathbf{x}) - m_{\theta 1}|] \stackrel{(i)}{\leq} \sqrt{\mathbb{E}_{\mathbf{x}} [(\hat{m}_1(\mathbf{x}) - m_{\theta 1})^2]} \stackrel{(ii)}{=} \frac{1}{\sqrt{\alpha}}, \quad (\text{B.5})$$

where (i) is by Jensen's inequality and equality (ii) holds as the variance of the Gaussian along the first coordinate is $1/\alpha$. Then, with this estimator in place

$$\begin{aligned} \Psi_2 &= \inf_{\mu} \frac{1}{|\Theta|} \sum_{\theta \in \Theta} \mathbb{E}_{\mathbf{x} \sim p_{\theta}} [\mathbb{E}_{\mathbf{t} \sim \mu_{\mathbf{x}}} [L_{\theta}(\mathbf{t})]] \stackrel{(i)}{=} \frac{1}{|\Theta|} \sum_{\theta \in \Theta} \mathbb{E}_{\mathbf{x} \sim p_{\theta}} [\min \{\alpha \|\hat{m}(\mathbf{x}) - m_{\theta}\|_2, 1\}] \\ &\leq \frac{1}{|\Theta|} \sum_{\theta \in \Theta} \alpha \mathbb{E}_{\mathbf{x} \sim p_{\theta}} [\|\hat{m}(\mathbf{x}) - m_{\theta}\|] \stackrel{(ii)}{\leq} \sqrt{\alpha}, \end{aligned}$$

where (i) follows by replacing the infimum over all estimator by our estimator and by the definition of L_{θ} , and (ii) follows by invoking Inequality (B.5) established above. This completes our proof. \square

Acknowledgements

We are grateful to Aditya Guntuboyina for pointing us towards the literature on Le Cam deficiency. We would also like to thank Jelena Diakonikolas, Sébastien Gerchinovitz, Michael Jordan, Aldo Pacchiano, Aaditya Ramdas and Morris Yau for many helpful conversations. We thank Kush Bhatia for helpful comments that improved the presentation of the results.

We gratefully acknowledge the support of the NSF through grants IIS-1619362 and IIS-1909365. Part of this work was completed while NC was interning at Google.

References

- [1] Alekh Agarwal, Peter Bartlett, Pradeep Ravikumar, and Martin Wainwright. Information-theoretic lower bounds on the oracle complexity of stochastic convex optimization. *IEEE Transactions on Information Theory*, 5(58):3235–3249, 2012.
- [2] Claude Bélisle, H. Edwin Romeijn, and Robert Smith. Hit-and-run algorithms for generating multivariate distributions. *Mathematics of Operations Research*, 18(2):255–266, 1993.
- [3] David Blackwell. Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, pages 265–272, 1953.
- [4] Henri Frédéric Bohnenblust, Lloyd Shapley, and Seymour Sherman. Reconnaissance in game theory. *RAND Corporation Research Memorandum Series*, 1949.
- [5] Nawaf Bou-Rabee and Martin Hairer. Nonasymptotic mixing of the MALA algorithm. *IMA Journal of Numerical Analysis*, 33(1):80–110, 2013.

- [6] Steve Brooks, Andrew Gelman, Galin Jones, and Xiao-Li Meng. *Handbook of Markov chain Monte Carlo*. CRC Press, 2011.
- [7] Sébastien Bubeck, Ronen Eldan, and Joseph Lehec. Sampling from a log-concave distribution with projected Langevin Monte Carlo. *Discrete & Computational Geometry*, 59(4):757–783, 2018.
- [8] Niladri S. Chatterji, Peter L. Bartlett, and Philip M. Long. Oracle lower bounds for stochastic gradient sampling algorithms. *arXiv preprint arXiv:2002.00291*, 2020.
- [9] Xiang Cheng and Peter Bartlett. Convergence of Langevin MCMC in KL-divergence. In *Proceedings of the Algorithmic Learning Theory*, 2018.
- [10] Xiang Cheng, Niladri Chatterji, Peter Bartlett, and Michael Jordan. Underdamped Langevin MCMC: A non-asymptotic analysis. In *Proceedings of the Conference On Learning Theory*, 2018.
- [11] Arnak Dalalyan. Theoretical guarantees for approximate sampling from smooth and log-concave densities. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 79(3):651–676, 2017.
- [12] Arnak Dalalyan and Avetik Karagulyan. User-friendly guarantees for the Langevin Monte Carlo with inaccurate gradient. *Stochastic Processes and their Applications*, 129(12):5278–5311, 2019.
- [13] Alain Durmus, Szymon Majewski, and Błażej Miasojedow. Analysis of Langevin Monte Carlo via convex optimization. *Journal of Machine Learning Research*, 20(73):1–46, 2019.
- [14] Alain Durmus and Eric Moulines. Nonasymptotic convergence analysis for the unadjusted Langevin algorithm. *The Annals of Applied Probability*, 27(3):1551–1587, 2017.
- [15] Raaz Dwivedi, Yuansi Chen, Martin Wainwright, and Bin Yu. Log-concave sampling: Metropolis-Hastings algorithms are fast. *Journal of Machine Learning Research*, 20(183):1–42, 2019.
- [16] Martin Dyer, Alan Frieze, and Ravi Kannan. A random polynomial-time algorithm for approximating the volume of convex bodies. *Journal of the ACM (JACM)*, 38(1):1–17, 1991.
- [17] Andreas Eberle, Arnaud Guillin, and Raphael Zimmer. Couplings and quantitative contraction rates for Langevin dynamics. *The Annals of Probability*, 47(4):1982–2010, 2019.
- [18] Donald Ermak. A computer simulation of charged particles in solution. I. Technique and equilibrium properties. *The Journal of Chemical Physics*, 62(10):4189–4196, 1975.
- [19] Rong Ge, Holden Lee, and Jianfeng Lu. Estimating normalizing constants for log-concave distributions: Algorithms and lower bounds. *arXiv preprint arXiv:1911.03043*, 2019.
- [20] Ulf Grenander and Michael Miller. Representations of knowledge in complex systems. *Journal of the Royal Statistical Society: Series B (Methodological)*, 56(4):549–581, 1994.
- [21] James Johndrow, Natesh Pillai, and Aaron Smith. No free lunch for approximate MCMC. *arXiv preprint arXiv:2010.12514*, 2020.

- [22] Olav Kallenberg. *Foundations of modern probability*. Springer Science & Business Media, 2006.
- [23] Ravi Kannan, László Lovász, and Miklós Simonovits. Isoperimetric problems for convex bodies and a localization lemma. *Discrete & Computational Geometry*, 13(3-4):541–559, 1995.
- [24] Lucien Le Cam. Sufficiency and approximate sufficiency. *The Annals of Mathematical Statistics*, pages 1419–1455, 1964.
- [25] Lucien Le Cam and Grace Lo Yang. *Asymptotics in statistics: Some basic concepts*. Springer Science & Business Media, 2012.
- [26] László Lovász. Hit-and-run mixes fast. *Mathematical Programming*, 86(3):443–461, 1999.
- [27] László Lovász and Miklós Simonovits. The mixing rate of Markov chains, an isoperimetric inequality, and computing the volume. In *Proceedings of the Symposium on Foundations of Computer Science*, pages 346–354. IEEE, 1990.
- [28] László Lovász and Miklós Simonovits. Random walks in a convex body and an improved volume algorithm. *Random structures & algorithms*, 4(4):359–412, 1993.
- [29] László Lovász and Santosh Vempala. Hit-and-run from a corner. *SIAM Journal on Computing*, 35(4):985–1005, 2006.
- [30] László Lovász and Santosh Vempala. The geometry of log-concave functions and sampling algorithms. *Random Structures & Algorithms*, 30(3):307–358, 2007.
- [31] Oren Mangoubi and Aaron Smith. Rapid mixing of Hamiltonian Monte Carlo on strongly log-concave distributions. *arXiv preprint arXiv:1708.07114*, 2017.
- [32] Oren Mangoubi and Nisheeth Vishnoi. Dimensionally tight bounds for second-order Hamiltonian Monte Carlo. In *Proceedings of the Advances in Neural Information Processing Systems*, 2018.
- [33] Kerrie Mengersen and Richard Tweedie. Rates of convergence of the Hastings and Metropolis algorithms. *The Annals of Statistics*, 24(1):101–121, 1996.
- [34] Radford Neal. MCMC using Hamiltonian dynamics. In *Handbook of Markov Chain Monte Carlo*, pages 139–188. Chapman and Hall/CRC, 2011.
- [35] Arkadii Nemirovsky and David Yudin. Problem complexity and method efficiency in optimization. *SIAM Review*, 27(2):264–265, 1985.
- [36] Giorgio Parisi. Correlation functions and computer simulations. *Nuclear Physics B*, 180(3):378–384, 1981.
- [37] Marcelo Pereyra. Proximal Markov Chain Monte Carlo algorithms. *Statistics and Computing*, 26(4):745–760, 2016.
- [38] Luis Rademacher and Santosh Vempala. Dispersion of mass and the complexity of randomized geometric algorithms. *Advances in Mathematics*, 219(3):1037–1069, 2008.
- [39] Maxim Raginsky and Alexander Rakhlin. Information-based complexity, feedback and dynamics in convex programming. *IEEE Transactions on Information Theory*, 57(10):7036–7056, 2011.
- [40] Aaditya Ramdas and Aarti Singh. Optimal rates for stochastic convex optimization under Tsybakov noise condition. In *Proceedings of the International Conference on Machine Learning*, 2013.

- [41] Christian Robert and George Casella. *Monte Carlo statistical methods*. Springer Science & Business Media, 2013.
- [42] Gareth Roberts and Osnat Stramer. Langevin diffusions and Metropolis-Hastings algorithms. *Methodology and computing in applied probability*, 4(4):337–357, 2002.
- [43] Gareth Roberts and Richard Tweedie. Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli*, 2(4):341–363, 1996.
- [44] Gareth Roberts and Richard Tweedie. Geometric convergence and central limit theorems for multidimensional Hastings and Metropolis algorithms. *Biometrika*, 83(1):95–110, 1996.
- [45] Robert Smith. Efficient Monte Carlo procedures for generating points uniformly distributed over bounded regions. *Operations Research*, 32(6):1296–1308, 1984.
- [46] Kunal Talwar. Computational separations between sampling and optimization. In *Proceedings of the Advances in Neural Information Processing Systems*, 2019.
- [47] Erik Torgersen. *Comparison of statistical experiments*. Cambridge University Press, 1991.
- [48] Martin Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.
- [49] Max Welling and Yee Teh. Bayesian learning via stochastic gradient Langevin dynamics. In *Proceedings of the International Conference on Machine Learning*, 2011.
- [50] Tatiana Xifara, Chris Sherlock, Samuel Livingstone, Simon Byrne, and Mark Girolami. Langevin diffusions and the Metropolis-adjusted Langevin algorithm. *Statistics & Probability Letters*, 91:14–19, 2014.